Subject: Re: [PATCH] new cgroup controller "fork"
Posted by Max Kellermann on Thu, 17 Feb 2011 14:09:25 GMT
View Forum Message <> Reply to Message

On 2011/02/17 14:50, KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com> wrote:
> I wonder allowing to set the limit to Root cgroup may imply the system death.
> How about disabling to set value to Root cgroup ?

That is taken care of already:

> > +static int
> > +cgroup_fork_populate(struct cgroup_subsys *ss, struct cgroup *cgroup)
> > +{
> > + if (cgroup->parent == NULL)
> > + /* cannot limit the root cgroup */
> > + return 0;

The attribute simply doesn't exist in the root cgroup.

Also watch the loop condition in cgroup_fork_pre_fork() closely, the
root cgroup isn't checked (even if you could find a way to configure
it):

> > + t = cgroup_fork_current();
> > + while (t->css.cgroup->parent != NULL && err == 0) {

> IIRC, fork()'s error code is EAGAIN or ENOMEM. The exisiting limit of
> rlimit() returns EAGAIN.
>
> How about -EAGAIN here ? I think it's not good to add new error code for
> system calls.

EPERM seemed appropriate to me, because the administrator disallows
more than N forks. If there are practical reasons for changing it to
EAGAIN or ENOMEM, I'm ok with that. Thanks for the hint.

Max

_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containe rs